arXiv:1611.07812v2 [cs.SE] 27 Nov 2016

Static Analysis of Communicating Processes using Symbolic Transducers (Extended version)

Vincent Botbol[1,2] vincent.botbol@cea.fr,
Emmanuel Chailloux[1] emmanuel.chailloux@lip6.fr, and
Tristan Le Gall[2] tristan.le_gall@cea.fr

[1]Sorbonne Universités, UPMC Univ Paris 06, CNRS, LIP6 UMR 7606, 4 place Jussieu 75005 Paris
[2]CEA, LIST, Software Reliability and Security Laboratory, 91191 Gif-sur-Yvette France

Abstract. We present a general model allowing static analysis based on abstract interpretation for systems of communicating processes. Our technique, inspired by Regular Model Checking, represents set of program states as lattice automata and programs semantics as symbolic transducers. This model can express dynamic creation/destruction of processes and communications. Using the abstract interpretation framework, we are able to provide a sound over-approximation of the reachability set of the system thus allowing us to prove safety properties. We implemented this method in a prototype that targets the MPI library for C programs.

This report is an extended version of the research paper submitted to VMCAI 2017.

1 Introduction

The static analysis of concurrent programs faces several well-known issues, including how to handle dynamical process creation. This last one is particularly challenging considering that the state space of the concurrent system may not be known nor bounded statically, which depends on the number and the type of variables of the program.

In order to overcome this issue, we combine a symbolic representation based on regular languages (like the one used in Regular Model Checking [1]) with a fixed-point analysis based on abstract interpretation [5]. We define the abstract semantics of a concurrent program by using of a symbolic finite-state transducer [15]. A (classical) finite-state transducer T encodes a set of rules to rewrite words over a finite alphabet. In a concurrent program, if each process only has a finite number of states, then we can represent a set of states of the concurrent program by a language and the transition function by a transducer. However, this assumption does not hold since we consider processes with infinite state space, so we have to represent a set of states of the concurrent program by a

lattice automaton [9] and its transition function by a lattice transducer, a new kind of symbolic transducers that we define in this paper. Lattice Automata are able to recognize languages over an infinite alphabet. This infinite alphabet is an abstract domain (intervals, convex polyhedra, etc.) that abstracts process states.

We show, on Fig. 1 (detailed in Sec. 2), the kind of programs our method is able to analyse. This program generates an unbounded sequence of processes $\{id = 0, x = 5\}; \{id = 1, x = 9\}; \{id = 2, x = 13\}; \ldots$ We want to prove safety properties such as: $x = 5 + \mathrm{id} \times 4$ holds for every process when it reached its final location l_9. The negation of this property is encoded as a lattice automaton *Bad* (Fig. 2) that recognizes the language of all bad configurations. Our verification algorithm is to compute an over-approximation of the reachability set *Reach*, also represented by a lattice automaton, then, by testing the emptiness of the intersection of the languages, we are able to prove this property : $\mathscr{L}(Reach) \cap \mathscr{L}(Bad) = \emptyset$.

```
1       if (id==0)
2         x := 1
3       else
4         receive(any_id,x);
5
6       create(next);
7       x := x+4;
8       send(next,x)
```

Fig. 1: Program example

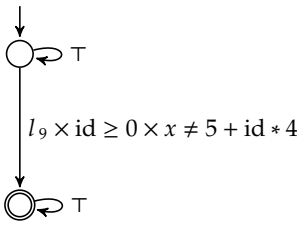

Fig. 2: Bad configurations

Related works. There are many works aiming at the static analysis of concurrent programs. Some of them use the abstract interpretation theory, but either they do not allow dynamic process creation [13] and/or use a different memory model [8] or do not consider numerical properties [7]. In [15], the authors defined symbolic transducers but they did not consider to raise it to the verification of concurrent programs. In [2], there is the same kind of representation that considers infinite state system but can only model finite-state processes. The authors of [4] present a modular static analysis framework targeting POSIX threads. Their model allows dynamic thread creation but lack communications between threads. More practically, [16] is a formal verification tool using a dynamical analysis based on model checking aiming at the detection of deadlocks in Message Passing Interface [14] (MPI) programs but this analysis is not sound and also does not compute the value of the variables.

Contributions. In this article, we define an expressive concurrency language with communication primitives and dynamic process creation. We introduce its concrete semantics in terms of symbolic rewriting rules. Then, we give a way to abstract multi-process program states as a lattice automaton and also abstract our semantics into a new kind of symbolic transducer and specific rules. We also give application algorithms to define a global transition function and prove

their soundness. A fixpoint computation is given to obtain the reachability set. Finally, in order to validate the approach, we implemented a prototype as a Frama-C [11] plug-in which targets a subset of MPI using the abstract domain library: Apron [10].

Outline. In Sec. 2, we present the concurrent language and its semantics definition, encoded by rewriting rules and a symbolic transducer. Then, Sec. 3 presents the abstract semantics and the algorithms used to compute the over-approximation of the reachability set of a program. In Sec. 4, we detail the implementation of our prototype targeting a subset of MPI which is mapped by the given semantics and run it on some examples (Sec. 5). We discuss about the potential and the future works of our method in Sec. 6.

2 Programming language and its Concrete Semantics

We present a small imperative language augmented with communications primitives such as unicast and multicast communications, and dynamical process creation. These primitives are the core of many parallel programming languages and libraries, such as MPI.

2.1 Language definition

In our model, memory is distributed: each process executes the same code, with its own set of variables. For the sake of clarity, all variables and expressions have the same type (integer), and we omit the declaration of the variables. Process identifiers are also integers.

⟨*program*⟩ ::= ⟨*instrs*⟩

⟨*instrs*⟩ ::= ⟨*instr*⟩ ';' ⟨*instrs*⟩

⟨*id*⟩ ::= ⟨*expr*⟩
| **any_id**

⟨*instr*⟩ ::= '{' ⟨*instrs*⟩ '}'
| ⟨*ident*⟩ ':=' ⟨*expr*⟩
| **if** '('⟨*expr*⟩')' ⟨*instr*⟩ [**else** ⟨*instr*⟩]
| **while** '('⟨*expr*⟩')' ⟨*instr*⟩
| **create** '(' ⟨*ident*⟩ ')'
| **send** '(' ⟨*id*⟩ ',' ⟨*ident*⟩ ')'
| **receive** '(' ⟨*id*⟩ ',' ⟨*ident*⟩ ')'
| **broadcast** '(' ⟨*expr*⟩ ',' ⟨*ident*⟩ ')'

⟨*ident*⟩ and ⟨*expr*⟩ stand for classical identifiers and arithmetic expressions on integers (as defined in the C language)

Communications are *synchronous*: a process with `id=orig` cannot execute the instruction `send(dest, var)` unless a process with `id=dest` is ready to execute the instruction `receive(orig, var')`; both processes then execute their instruction and the value of `var` (of process `orig`) is copied to variable `var'` (of process `dest`). We also allow unconditional receptions with `all_id` meaning that a process with `id=orig` can receive a variable whenever another process is ready

to execute an instruction `send(orig,v)`. `broadcast(orig, var)` instructions cannot be executed unless all processes reach the same instruction. `create(var)` dynamically creates a new process that starts its execution at the program entry point. The id of the new process, which is a fresh id, is stored in `var`, so the current process can communicate with the newly created process. Other instructions are asynchronous. Affectations, conditions and loops keep the same meaning as in the *C* language.

2.2 Formal Semantics

We model our program using an unbounded set P of processes, ordered by their identifiers ranging from 1 to $|P|$. As usual, the control flow graph (CFG) of the program is a graph where vertices belong to a set L of program points and edges are labelled by a $instr \subseteq L \times Instr \times L$ where $Instr$ are the instructions defined in our language. Finally, V represents the set of variables. Their domain of values is $\mathbb{V} \supseteq \mathbb{N}$. For any expression $expr$ of our language, and any valuation $\rho : V \to \mathbb{V}$, we note $\text{eval}(expr, \rho) \in \mathbb{V}$ its value.

Our processes share the same code and have distributed memory: each variable has a local usage in each process. Thus, a *local state* is defined as $\sigma \in \Sigma = Id \times L \times (V \to \mathbb{V})$. It records the identifier of the process, its current location and the value of each local variable.

A *global state* is defined as a *word* of process local states: $\sigma_1 \cdot \sigma_2 \cdot ... \cdot \sigma_n \in \Sigma^*$ where n is the number of running processes and Σ^* is the free monoid on Σ.

The semantics is given as a transition system $\langle \Sigma^*, I, \tau \rangle$, where $I \in \Sigma^*$ is the set of all possible initial program states. As the code is shared, every process starts at the same location l_0 and every variable's value is initialised with 0. Therefore, if there are initially n processes, $I = \{\sigma_1 \cdot ... \cdot \sigma_n\}$ where $\forall i \in n, \sigma_i = \langle i, l_0, (\lambda v\ .\ 0) \rangle$. The transition relation $\tau \subseteq \Sigma^* \times \Sigma^*$ is defined as:

- for each local instruction (e.g. assignments, conditionals, and loops) $(l, a, l') \in instr$, we have:
  $((\sigma_1 \cdot ... \cdot \langle id, l, \rho \rangle \cdot ... \cdot \sigma_n), (\sigma_1 \cdot ... \cdot \langle id, l', \rho' \rangle \cdot ... \cdot \sigma_n)) \in \tau$
  where $\rho' = [\![a]\!]\rho$ is the classical small-step semantics of action a
- for every pair of `send/receive` instructions of two processes :
  $(l_i, send(id\_to, v_i), l'_i)$ and $(l_j, receive(id\_from, v_j), l'_j)$ (or $(l_j, receive(\texttt{id\_all}, v_j), l'_j)$),
  we have:
  $((\sigma_1 \cdot ... \cdot \langle id_i, l_i, \rho_i \rangle \cdot ... \cdot \langle id_j, l_j, \rho_j \rangle \cdot ... \cdot \sigma_n), (\sigma_1 \cdot ... \cdot \langle id_i, l'_i, \rho_i \rangle \cdot ... \cdot \langle id_j, l'_j, \rho'_j \rangle \cdot ... \cdot \sigma_n)) \in \tau$
  where $\rho'_j = \rho_j[v_j \leftarrow \rho_i(v_i)]$ when $(id\_from = \texttt{id\_all}$ or $id_i = \text{eval}(id\_from, \rho_j))$
  and $(id\_to = \texttt{id\_all}$ or $id_j = \text{eval}(id\_to, \rho_i)$)
- for each `broadcast` instruction $(l, broadcast(id_x, v), l')$
  $((\langle id_1, l, \rho_1 \rangle \cdot ... \cdot \langle id_x, l, \rho_x \rangle \cdot ... \cdot \langle id_n, l, \rho_n \rangle), (\langle id_1, l', \rho'_1 \rangle \cdot ... \cdot \langle id_x, l', \rho_x \rangle \cdot ... \cdot \langle id_n, l', \rho'_n \rangle)) \in \tau$
  where $\forall i \in [1, n], id_i \neq id_x \Rightarrow \rho'_i = \rho_i[v \leftarrow \rho_x(v)]$

– finally, for each `create` instruction $(l, create(v), l')$
$((\sigma_1 \cdot ... \cdot \langle id, l, \rho\rangle \cdot ... \cdot \sigma_n), (\sigma_1 \cdot ... \cdot \langle id, l', \rho'\rangle \cdot ... \cdot \sigma_n \cdot \sigma_{n+1})) \in \tau$
where $\rho' = \rho[v \leftarrow n+1]$ and $\sigma_{n+1} = \langle n+1, l_0, (\lambda v\ .\ 0)\rangle$

In the following, we directly consider sets $E \in \mathscr{P}(\Sigma^*)$ and Post_τ defined as:

$$\text{Post}_\tau(E) = \{w' \in \Sigma^* \mid \exists w \in E\ \wedge\ (w, w') \in \tau\}$$

Post^*_τ is the reflexive and transitive closure of Post_τ. Given an initial set of states $I \in \mathscr{P}(\Sigma^*)$, the *reachability set* $\text{Post}^*_\tau(I)$ contains all states that can be found during an execution of the program. Assuming we want to check whether the program satisfies a safety property (expressed as a *bad configuration*) given by a set of states B that must be avoided, the verification algorithm is simply to test whether $\text{Post}^*_\tau(I) \cap B = \emptyset$; if true, the program is safe.

Therefore, we would like to define Post_τ in a more *operational* way, as a set of rewriting rules that can be applied to I, so we can apply those rules iteratively until we reach the fixpoint $\text{Post}^*_\tau(I)$.

2.3 Symbolic Rewriting Rules

Let us consider a local instruction $(l, a, l') \in instr$; for any set of states E:

$$\text{Post}_{(l,a,l')}(E) = \{\sigma_1 \cdot ... \cdot \langle id, l', \rho'\rangle \cdot ... \cdot \sigma_n) \mid \exists \sigma_1 \cdot ... \cdot \langle id, l, \rho\rangle \cdot ... \cdot \sigma_n) \in E\ \wedge\ \rho' = [\![a]\!]\rho\}$$

The effects of $\text{Post}_{(l,al')}$ on E is to rewrite every word of E. Thus, we would like to express it as a rewriting rule G/F where G is a symbolic guard matching a set of words and F a symbolic rewriting function. Since our method uses the framework of abstract interpretation (see Sec. 3), *symbolic* means that we consider elements of some lattice to define the rules. We give the rewriting rule that encodes the execution of a local instruction (l, a, l'):

$$G = \top^* \cdot \langle \_, l, \_\rangle \cdot \top^* \text{ and } F = \texttt{Id}^* \cdot f \cdot \texttt{Id}^* \text{ with } f(X) = \{\langle id, l', [\![a]\!](\rho)\rangle | \langle id, l, \rho\rangle \in X\}$$

The guard matches words composed of any number of processes, then one process with location l, then again any number of processes. The function $\texttt{Id}^*$ means that the processes matched by $\top^*$ will be rewritten as the identity and therefore not modified. $\Lambda = \mathscr{P}(\Sigma)$ is the lattice of sets of local states. $f : \Lambda \to \Lambda$ rewrites a set of local states according to the semantics of a. So every word $w \in E$ that matches the guard will be rewritten and we will obtain $\text{Post}_{(l,a,l')}(E)$.

We now give the general definition of those rewriting rules and how to apply them. We remind that the partial order $\sqsubseteq$ can be extended to Λ^* as $u \sqsubseteq v$ if both words have the same length ($|u| = |v|$) and $\forall i < |u|, \bot \neq u_i \sqsubseteq v_i$. Note that we do not allow $\bot$ in words: any word that would contain one or more $\bot$ letters is identified to the smallest element $\bot_{\Lambda^*}$. Therefore, any word $w \in \Lambda^*$ represents a set of words of Σ^*: $\sigma_1 \ldots \sigma_n \in w$ when $\{\sigma_1\} \ldots \{\sigma_n\} \sqsubseteq w$.

Definition 1 *Let Λ be a lattice. A rewriting rule over Λ is given by two sequences $G = (g_0)^* \cdot w_1 \cdot (g_1)^* \cdot w_2 \dots w_n \cdot (g_n)^*$ and $F = f_0 \cdot h_0 \cdot f_1 \cdot h_1 \cdot ... \cdot h_n \cdot f_{n+1}$ such that:*

- $\forall\ 1 \leq i \leq n, w_i \in \Lambda^*$ *and* $|w_i| > 0$;
- $\forall\ 0 \leq i \leq n, g_i \in \Lambda$;
  We note $N = |w_1| + |w_2| + ... + |w_n|$
- $\forall\ 0 \leq i \leq n+1, f_i : \Lambda^N \to \Lambda^*$;
- $\forall\ 0 \leq i \leq n, h_i : \Lambda^{N+1} \to \Lambda$.

With this rule, a finite word $w \in \Lambda^*$ is rewritten to $w' \in \Lambda^*$ if:

- *w can be written as a concatenation $w = u_0 \cdot v_1 \cdot u_1 \cdot ... \cdot v_n \cdot u_n$ with:*
  - $\forall\ 0 \leq i \leq n, u_i = \lambda_0 \dots \lambda_{|u_i|}$ *and* $\forall\ 0 \leq j \leq |u_i|, \lambda_j \sqsubseteq g_i$,
  - $\forall\ 1 \leq i \leq n, v_i \sqsubseteq w_i$;
- *$w' = v'_0 \cdot u'_0 \cdot v'_1 \cdot u'_1 \cdot ... \cdot v'_n \cdot u'_n \cdot v'_{n+1}$ with:*
  - $\forall\ 0 \leq i \leq n, u'_i = \lambda'_0 \cdot \lambda'_1 \cdot ... \cdot \lambda'_{|u_i|}$ *and* $\forall\ 0 \leq j \leq |u_i|, \lambda'_j = h_i(\lambda_j, v_1, \dots, v_n)$,
  - $\forall\ 0 \leq i \leq n+1, v'_i = f_i(v_1, \dots, v_n)$.

For any $N \in \mathbb{N}$, $\texttt{Id}^* : \Lambda^{N+1} \to \Lambda$ is defined as $\texttt{Id}^*(x, y_1, \dots y_N) = x$. Moreover, we denote by $\langle \_, l, \_ \rangle$ the element of $\Lambda = \mathscr{P}(\Sigma)$ defined as $\{\langle id, l, \rho \rangle \mid \forall id, \rho\}$, (the symbol '_' matches anything). With these notations, we can express the transition relation by a set of rewriting rules:

- For every pair of `send/receive` instructions $(l_i, \texttt{send(id\_to, v}_\texttt{i}\texttt{)}, l'_i)$ and $(l_j, \texttt{receive(id\_from, v}_\texttt{j}\texttt{)}, l'_j)$, we have the rule:

$$G = \top^* \cdot \langle \_, l_i, \_ \rangle \cdot \top^* \cdot \langle \_, l_j, \_ \rangle \cdot \top^* \text{ and } F = \texttt{Id}^* \cdot f_1 \cdot \texttt{Id}^* \cdot f_2 \cdot \texttt{Id}^* \text{ with}$$

  - $f_1(E_1, E_2) = \{\langle id_i, l'_i, \rho_i \rangle \mid \langle id_i, l_i, \rho_i \rangle \in E_1 \wedge \langle id_j, l_j, \rho_j \rangle \in E_2 \wedge id_i = \text{eval}(id\_from, \rho_j) \wedge id_j = \text{eval}(id\_to, \rho_i)\}$
  - $f_2(E_1, E_2) = \{\langle id_j, l'_j, \rho_j[v_j \leftarrow \text{eval}(v_i, \rho_i)] \rangle \mid \langle id_i, l_i, \rho_i \rangle \in E_1 \wedge \langle id_j, l_j, \rho_j \rangle \in E_2 \wedge id_i = \text{eval}(id\_from, \rho_j) \wedge id_j = \text{eval}(id\_to, \rho_i)\}$

  and symmetrically when σ_j is located before σ_i in the word of local states. When e.g. $id\_to = \texttt{id\_all}$, the condition $id_j = \text{eval}(id\_to, \rho_i)$ is satisfied for any (id_j, ρ_j).
- for each `broadcast` instruction $(l, \texttt{broadcast(id\_x, v)}, l')$, we have the rule:

$$G = (\langle \_, l, \_ \rangle)^* \cdot \langle id\_x, l, \_ \rangle \cdot (\langle \_, l, \_ \rangle)^* \text{ and } F = F_1^* \cdot f_1 \cdot F_1^* \text{ with}$$

  - $F_1^*(E_1, E_2) = \{\langle id_i, l', \rho_i[v \leftarrow \text{eval}(v, \rho_x)] \rangle \mid \langle id_i, l, \rho_i \rangle \in E_1 \wedge \langle id_x, l, \rho_x \rangle \in E_2\}$
  - $f_1(E_1) = \{\langle id_x, l', \rho_x \rangle \mid \langle id_x, l, \rho_x \rangle \in E_1\}$

  The guard $\langle id\_x, l, \_ \rangle$ stands for the set $\{\langle id_i, l_i, \rho_i \rangle \mid l_i = l \wedge id_i = \text{eval}(id\_x, \rho_i)\}$
- finally, for each `create` instruction $(l, \texttt{create(v)}, l')$, we have the rule:

$$G = \top^* \cdot \{\_, l, \_\} \cdot \top^* \text{ and } F = Id^* \cdot f_1 \cdot Id^* \cdot f_2 \text{ with}$$

  - $f_1(E_1) = \{\langle id_i, l', \rho_i \rangle \mid \langle id_i, l, \rho_i \rangle \in E_1\}$
  - $f_2(E_1) = \{\langle id_n, l_0, (\lambda v\ .\ 0) \rangle \mid n = \texttt{fresh\_id}()\}$

where `fresh_id` returns a new unique identifier n where $n = |w| + 1$ with w the word of processes.

Example 1. Let us consider consider our running example depicted on Fig. 1. Let us assume we have a set of program states $E = \{\langle id = 0, l_0, x = 0, next = 0\rangle; \langle id = 0, l_9, x = 5, next = 1\rangle.\langle id = 1, l_8, x = 9, next = 2\rangle.\langle id = 2, l_4, x = 0, next = 2\rangle; \langle id = 1, l_8, x = 13, next = 2\rangle.\langle id = 6, l_4, x = 0, next = 0\rangle.\}$, i.e. there is either one process in l_0, or three process in l_9, l_8, l_4 or two processes in l_4, l_8. We consider the symbolic rewriting rule that results from the communication instructions. Its guard is $\top^*.\langle \_, l_8, \_\rangle.\top^*.\langle \_, l_4, \_\rangle.\top^*$ and its rewriting functions $\mathtt{Id}^* \cdot f_1 \cdot \mathtt{Id}^* \cdot f_2 \cdot \mathtt{Id}^*$ with

- $f_1(E_1, E_2) = \{\langle id_i, l_9, \rho_i\rangle \mid \langle id_i, l_8, \rho_i\rangle \in E_1 \wedge \langle id_j, l_4, \rho_j\rangle \in E_2 \wedge id_j = \text{eval}(next, \rho_i)\}$
- $f_2(E_1, E_2) = \{\langle id_j, l_5, \rho_j[x \leftarrow \text{eval}(x, \rho_i)]\rangle \mid \langle id_i, l_8, \rho_i\rangle \in E_1 \wedge \langle id_j, l_4, \rho_j\rangle \in E_2 \wedge id_j = \text{eval}(next, \rho_i)\}$

then $\text{Post}_\tau(E) = \{\langle id = 0, l_9, x = 9, next = 1\rangle.\langle id = 1, l_9, x = 13, next = 2\rangle.\langle id = 2, l_5, x = 13, next = 2\rangle\}$, which is the image of the state with three active processes. There is no possible communication when $\langle id = 1, l_8, x = 13, next = 2\rangle.\langle id = 6, l_4, x = 0, next = 0\rangle$. Even if the locations match the guard, the first process can only send messages to a process with $\mathtt{id} = \text{eval}(next, \rho) = 2 \neq 6$.

Transducers Alternatively, the semantics of local instructions can also be described by a *lattice transducer*. A finite-state transducer is a finite-state automaton but instead of only accepting a language, it also rewrites it. A lattice transducer is similar to a finite-state transducer; however, it is symbolic, i.e. it accepts inputs (and produces outputs) belonging to the lattice Λ, which may be an infinite set.

Definition 2 *A* Lattice Transducer *is a tuple* $T = \langle \Lambda, Q, Q_0, Q_f, \Delta\rangle$ *where:*

- Λ *is a lattice*
- Q *is a finite set of states*
- $Q_0 \subseteq Q$ *are the initial states set*
- $Q_f \subseteq Q$ *are the final states set*
- $\Delta \subseteq Q \times \Lambda^n \times (\Lambda^n \to \Lambda)^* \times Q$ *with* $n \in \mathbb{N}^0$ *is a finite set of transitions with guards and rewriting functions*

Let $w = \lambda_1 \cdot \ldots \cdot \lambda_n \in \Lambda^n$ and $\{q, G, F, q'\} \in \Delta$ with $G = \gamma_1, ..., \gamma_n$ and $F = f_1, ..., f_m$. We write $q \xrightarrow{w/w'} q'$ when:

$$\begin{cases} \forall i \in [1, n] \;\; \lambda_i \sqsubseteq \gamma_i \\ w' = f_1(\lambda_1, ..., \lambda_n) \cdot \ldots \cdot f_m(\lambda_1, ..., \lambda_n) \end{cases}$$

For any word $w \in \Lambda^*$, $T(w)$ is the set of words w' such that there exists a sequence $q_0 \xrightarrow{w_1/w'_1} q_1 \xrightarrow{w_2/w'_2} \ldots \xrightarrow{w_n/w'_n} q_f$ with $q_0 \in Q_0$, $q_f \in Q_f$, $w = w_1.w_2 \ldots w_n$ and $w' = w'_1.w'_2 \ldots w'_n$. For any language $L \subseteq \Sigma^*$, $T(L) = \cup_{w \in L} T(w)$.

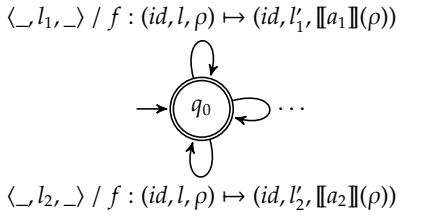

Fig. 3: Local transitions Transducer

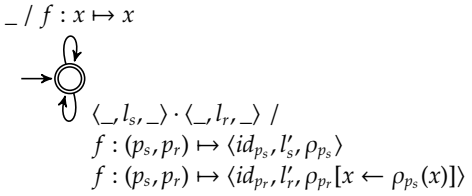

Fig. 4: "Neighbour" communication

We can express the semantics of the local instructions $\langle l_1, a_1, l'_1 \rangle, \langle l_2, a_2, l'_2 \rangle, \ldots$ by a transducer as shown in Fig. 3.

For the language we presented, the transducer representation is not fully exploited. Indeed, only single self-looping transitions are present. Yet, in our example program, we notice that communications and dynamic creation are done in their "neighbourhood": processes send their x to their right neighbor, receive from the left and create processes on their right-side. This semantics can be expressed with our transducer representation. We give on Fig. 4 a transducer encoding a "neighbour" version of synchronous communications as `send_right` and `receive_left` primitives. In our illustration, we use the locations (l_s, l'_s) and (l_r, l'_r) in order to represent pre and post locations of `send_right` and `receive_left` instructions. However, this restriction is not satisfying: we wish to handle point-to-point communications regardless of process locations in words of states. Thus we have to limit the transducer to encode only local transitions.

Therefore, communications are encoded by semantics rules R, and local instructions by a transducer T. We note T_{ext} the transducer extended with semantic rules, i.e. for any language $X \subseteq \Sigma^*$, $T_{ext}(X) = R(X) \cup T(X) = \text{Post}_\tau(X)$. For any initial set of states $I \subseteq \{\mathscr{P}(\Sigma)\}$, we have the reachability set $\text{Post}^*_\tau(I) = T^*_{ext}(I)$. However, $T^*_{ext}(I)$ cannot be computed in general, so we need abstractions.

3 Abstract Semantics

3.1 Lattice Automata

We give here a look at the lattice automata. The reader may refer to [9] for further details. As said before, the definition of lattice automata requires Λ to be atomistic, i.e.:

- Atoms(Λ) is the set of *atoms*; $\lambda \in \Lambda$ is an atom if $\forall \lambda' \in \Lambda, \lambda' \sqsubseteq \lambda \;\Rightarrow\; \lambda' = \lambda \vee \lambda' = \bot$
- Λ is atomic, i.e. : $\forall \lambda \in \Lambda,\ \lambda \neq \bot \Rightarrow \exists \lambda' \in \text{Atoms}(\Lambda), \lambda' \sqsubseteq \lambda$
- any element is equal to to least upper bound of atoms smaller than itself: $\forall \lambda \in \Lambda,\ \lambda = \sqcup\{\lambda' \mid \lambda' \in \text{Atoms}(\Lambda), \lambda' \sqsubseteq \lambda\}$

The language recognized by lattice automata are on the set of atoms rather than on Λ itself. The reason for this is that there may be different edges between the two same nodes. For example, let us consider the lattice of intervals, and let us consider the three automata depicted on Fig. 5. Intuitively, they represent the same set, but if we define their language as: $w \in \Lambda^*, q_0 \overset{w}{\twoheadrightarrow} q_f$, $\mathscr{L}(A_1) = \{[0,2];[2,4]\}$ while $\mathscr{L}(A_2) = \{[0,3];[3,4]\}$. If we define the language on atoms, both automata recognize the language: $\{[0,0];[1,1];[2,2];[3,3];[4,4]\}$ (assuming we only consider integer bounds). We can also merge transitions and have automaton A_3 that recognizes the same language.

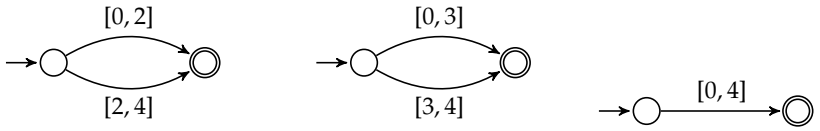

(a) A_1 (b) A_2 (c) A_3

Fig. 5: Three equivalent lattice automata

Thus the definition of the language allow us to split or merge transitions as long as the language remain the same. But if the interval $[0, +\infty]$ may be split in an infinite number of smaller intervals, how can we ensure that there is only a finite number of transitions ? We introduce an arbitrary, finite partition π of the atoms. π may be defined as a function $\pi : K \to \Lambda$, where K is an arbitrary finite set, such that if $k_1 \neq k_2$, $\pi(k_1) \sqcap \pi(k_2) = \bot$ and $\forall a \in \text{Atoms}(\Lambda), \exists k \in K, a \sqsubseteq \pi(k)$.

We define *Partitioned Lattice Automata* (PLAs) as the automata such that for any transition $(p, \lambda, q) \in \Delta_A, \exists k \in K, \lambda \sqsubseteq \pi(k)$ (i.e. all the atoms smaller than λ belong to the same partition class). A PLA is *merged* if $(p, \lambda, q) \in \Delta_A \wedge (p, \lambda', q) \in \Delta_A \Rightarrow \pi^{-1}(\lambda_1) \neq \pi^{-1}(\lambda_2)$, i.e. there is at most one transition per element of the partition. So merged PLAs have a finite number of transitions. Moreover, we can use this partition to design algorithms similar to the ones for Finite State Automata (such as union, intersection, determinisation and minimisation), with K playing the role of a finite alphabet. Indeed, if A is a merged PLA, we can apply π^{-1} to every label of the transitions and obtain a finite-state automata called $shape(A)$. *Normalised PLAs* are merged PLAs that are also deterministic and minimised.

If we have ∇_{auto}, a widening operator on finite-state automata, and ∇_Λ a widening operator on Λ then we have a widening operator on lattice automata $A_1 \nabla A_2$:

- if $shape(A_1)$ and $shape(A_2)$ are isomorphic, then we apply ∇_Λ on pairs of isomorphic transitions
- otherwise we compute $shape(A_1) \nabla_{auto} shape(A_2)$ and then merge transitions accordingly.

Therefore, lattice automata are a convienient way to "lift" a numerical domain Λ to an abstract domain for languages over Atoms(Λ), and to extend static analysis of sequential programs to concurrent programs. They can also easily handle disjunctive local invariants: $\lambda_1 \vee \lambda_2$ is simply represented by two transitions (p, λ_1, q) and (p, λ_2, q). Moreover, the whole reachability set is represented by a single automaton, which is both a blessing and a curse: it provides a concise, graphical way to represent the rechability set, but it also means that when computing a fixpoint by iteration (e.g. computing $T^*(A)$), we compute an increasing sequence of (increasingly large) automata $A_{i+1} = A_i \cup T(A_i)$. When applying T to A_{i+1}, we have $T(A_{i+1}) = T(A_i) \cup T(T(A_i))$ should avoid to recompute $T(A_i)$ (either using cache or having a way to apply T only to the 'increment').

3.2 Lattice Automata as an abstract domain

Since Σ may be an infinite set, we must have a way to abstract languages (i.e. subsets of Σ^*) over an infinite alphabet. *Lattice Automata* [9] provide this kind of abstractions. Lattice Automata are similar to finite-state automata, but their transitions are labeled by elements of a lattice. In our case, lattice automata are appropriate because:

- they provide a finite representation of languages over an infinite alphabet;
- we can apply symbolic rewriting rules or a transducer to a lattice automaton (see Sec. 3.3);
- there is a widening operator that ensures the termination of the analysis (see Sec. 3.4).

Definition 3 *A lattice automaton is defined by a tuple $A = \langle \Lambda, Q, Q_0, Q_f, \delta \rangle$ where:*

- *Λ is an atomistic lattice*[1]*, the order of which is denoted by $\sqsubseteq$;*
- *Q is a finite set of states;*
- *$Q_0 \subseteq Q$ and $Q_f \subseteq Q$ are the sets of initial and final states;*
- *$\Delta \subseteq Q \times (\Lambda \setminus \{\bot\}) \times Q$ is a finite transition relation.*[2]

This definition requires Λ to have a set of atoms Atoms(Λ). Abstract lattices like Intervals [5], Octagons [12] and Convex Polyhedra [6] are atomistic, so we can easily find such lattices to do our static analysis. Note that if Λ is atomistic, Λ^N and Λ^* are also atomistic, their atoms belonging to respectively Atoms(Λ)^N and Atoms(Λ)*. Moreover, for any set Σ, the lattice $\mathscr{P}(\Sigma, \subseteq)$ is atomistic and its atoms are the singletons. In the remainder of this paper, we will assume that any lattice we consider is atomistic. Finally, in addition to a widening operator, lattice automata have classic FSA operations ($\cup$, $\cap$, $\subseteq$, etc.).

The language recognized by a lattice automaton A is noted $\mathscr{L}(A)$ and is defined by finite words on the alphabet Atoms(Λ). $w \in \mathscr{L}(A)$ if $w = \lambda_1 \dots \lambda_n \in$

[1] See [9] or Appendix 3.1.

[2] No transition is labeled by $\bot$.

Atoms$(\Lambda)^*$ and there is a sequence of states and transitions $q_0 \xrightarrow{\lambda_0} q_1 \xrightarrow{\lambda_2} \ldots \xrightarrow{\lambda_n} q_n$ with $q_0 \in Q_0$ and $q_n \in Q_f$.

The reason why we define the language recognized by a lattice automaton as sequence of atoms are discussed in [9]; in a nutshell, this definition implies that two lattice automata that have the same concretisation recognize the same language. Moreover, by introducing a finite partition of the atoms, we can define determinisation and minimisation algorithms similar to the ones for finite-state automata, as well as a canonical form (*normalized* lattice automata).

Abstractions and Concretisations Assuming there is a Galois connection between $\mathscr{P}(\Sigma)$ and Λ we can extend the concretisation function $\gamma : \Lambda \to \mathscr{P}(\Sigma)$, we can extend it to $\gamma : \Lambda^* \to \mathscr{P}(\Sigma^*)$; if $w = \lambda_1 \ldots \lambda_n \in \Lambda^*$, $\gamma(w) = \{\sigma_1 \ldots \sigma_n | \forall i = 1..n, \sigma_i \in \gamma(\lambda_i)\}$ and for any language L, $\gamma(L) = \cup_{w \in L} \gamma(w)$. Thus, the concretisation of a lattice automaton A is $\gamma(\mathscr{L}(A))$, which can be computed by applying γ to all of A. Lattice automata are not a complete lattice; the abstraction function is defined as: if L is regular (i.e. it can be represented by a lattice automaton with labels in $\mathscr{P}(\Sigma)$) $\alpha(L)$ we can apply α to each edge; otherwise $\alpha(L) = \top$. The latter case does not happen in practice, since the initial set of states I is regular, and since we only check regular properties. We now present algorithms to apply a symbolic rewriting rule or a lattice transducer to a lattice automaton.

3.3 Algorithms

Application of a Rule To apply a symbolic rewriting rule to the language recognized by a lattice automaton, we must first identify the subset of words that match the guard $(g_0)^* \cdot w_1 \cdot (g_1)^* \cdot w_2 \ldots w_n \cdot (g_n)^*$. In this guard, it's easier to look first for sequences in the automaton that match $w_1, w_2, \ldots, w_n$. In automaton A a sequence that matches e.g. w_1 begins from state q_b^1 and ends in state q_e^1. Then, we identify the sub-automaton that could match $(g_0)^*$, i.e. all the states that are reachable from an initial state q_0 and correachable from q_b^1 by considering only transitions labeled by elements λ such that $g_0 \sqcap \lambda \neq \bot$. Once each part is identified, we can apply the rewriting function to each part and then we get a new automaton A'. Since this pattern matching is non deterministic, we have to consider all possible matching sequences. The result of the algorithm is the union of every automaton A' constructed in this way.

We introduce some notations before writing the algorithm. Let $w = \lambda_1 \ldots \lambda_n \in \Lambda^n$ and let A be a lattice automaton. We denote by matches(w, A) the set of matching sequences:

$$\begin{aligned}\text{matches}(w, A) = \{(q_b, v_1 \ldots v_n, q_e) \mid \exists q_0 \xrightarrow{\lambda'_1} q_1 \xrightarrow{\lambda'_2} \ldots \xrightarrow{\lambda'_n} q_n \in A,\\ q_0 = q_b \wedge q_n = q_e \wedge \forall i = 1..n, v_i = \lambda_i \sqcap \lambda'_i \neq \bot\}\end{aligned}$$

Let (q_b, q_e) be a pair of states of a lattice automaton $A = \langle \Lambda, Q, Q_0, Q_f, \delta \rangle$ and let $\lambda \in \Lambda$. We denote by $A_{q_b \to q_e}$ the sub-automaton $A_{q_b \to q_e} = \langle \Lambda, Q, \{q_b\}, \{q_e\}, \delta \rangle$.

For a lattice automaton $A = \langle \Lambda, Q, Q_0, Q_f, \delta \rangle$ and a function $f : \Lambda \to \Lambda$, we denote by $\text{map}(f, A)$ the automaton $\text{map}(f, A) = \langle \Lambda, Q, Q_0, Q_f, f(\delta) \rangle$ where $f(\delta) = \{(q, f(\lambda), q') | (q, \lambda, q') \in \delta \wedge f(\lambda) \neq \bot\}$.

With those notations, we give an algorithm to apply a rewriting rule on a lattice automaton:

ApplyRule ($G = (g_0)^* \cdot w_1 \cdot (g_1)^* \ldots w_n \cdot (g_n)^*$ and $F = f_0 \cdot h_0 \cdot f_1 \cdot ... \cdot h_n \cdot f_{n+1}$, A) :
Result := $\emptyset$
For all matching sequences
$(q_b^1, v^1, q_e^1) \in \text{matches}(w_1, A)$, ... , $(q_b^n, v^n, q_e^n) \in \text{matches}(w_n, A)$,
for each initial state $q_0 \in Q_0^A$ and for each final state $q_f \in Q_f^A$
(Let $A_0 = \text{map}(x \mapsto g_0 \sqcap x, A)_{q_0 \to q_b^1}$,
$A_1 = \text{map}(x \mapsto g_1 \sqcap x, A)_{q_e^1 \to q_b^2}$, ... , $A_n = \text{map}(x \mapsto g_n \sqcap x, A)_{q_e^n \to q_f}$.
For $i = 0 \,..\, n$:
let $A'_i = \text{map}(x \mapsto h_i(x, v^1, \ldots, v^n), A_i)$.
For $i = 0 \,..\, n+1$:
let $w'_i = f_i(v^1, \ldots, v^n)$.
Let q_{-1} and q_{n+1} be two fresh states (not appearing in any A'_i).
Let $\delta^{seq} = \{(q_{-1}, w'_0, q_0)(q_b^1, w'_1, q_e^1)(q_b^1, w'_1, q_e^1) \ldots (q_b^n, w'_n, q_e^n)(q_b^n, w'_{n+1}, q_{n+1})\}$.
Let $A' = \langle \Lambda, Q \cup \{q_{-1}, q_{n+1}\}, \{q_{-1}\}, \{q_{n+1}\}, \delta^{A'} \rangle$ with
$\delta^{A'} = \delta^{seq} \cup \delta^{A'_0} \cup \cdots \cup \delta^{A'_n}$

Result := Result $\cup A'$)
return Result

Theorem 1. *Let $R = (g_0)^* \cdot w_1 \cdot (g_1)^* \cdot w_2 \ldots w_n \cdot (g_n)^*$ / $f_0 \cdot h_0 \cdot f_1 \cdot h_1 \cdot ... \cdot h_n \cdot f_{n+1}$ be a rewriting rule and A a lattice automaton. If $R(A) = \text{ApplyRule}(R, A)$, then we have: $R(\mathscr{L}(A)) \subseteq \mathscr{L}(R(A))$.*

Proof. Let $R = (g_0)^* \cdot w_1 \cdot (g_1)^* \cdot w_2 \ldots w_n \cdot (g_n)^* / f_0 \cdot h_0 \cdot f_1 \cdot ... \cdot h_n \cdot f_{n+1}$

Let $w \in \mathscr{L}(A)$. If w matches the guard $(g_0)^* \cdot w_1 \cdot (g_1)^* \cdot w_2 \ldots w_n \cdot (g_n)^*$, it means we can decompose it as $w = u_0.v_1.u_1.v_2.u_2 \ldots v_n.u_n$ such that:

- for $i = 1..n$, $v_i \sqsubseteq w_i$
- for $i = 0..n$, each letter of u_i is smaller than g_i

Since $w \in \mathscr{L}(A)$, we consider a path $q_0 \xrightarrow{u_0}\!\!\!\!\rightarrow q_b^1 \xrightarrow{v_1}\!\!\!\!\rightarrow q_e^1 \ldots \xrightarrow{v_n}\!\!\!\!\rightarrow q_e^n \xrightarrow{u_n}\!\!\!\!\rightarrow q_f$ in A, i.e. there are matching sequences $(q_b^1, v^1, q_e^1) \in \text{matches}(w_1, A)$, $(q_b^2, v^2, q_e^2) \in \text{matches}(w_2, A)$, ... , $(q_b^n, v^n, q_e^n) \in \text{matches}(w_n, A)$, such that $\forall i, v_i \sqsubseteq v^i$. In algorithm `ApplyRule`, these matching sequences generate an automaton A'. By applying the rewriting functions $f_0 \cdot h_0 \cdot f_1 \cdot h_1 \cdot ... \cdot h_n \cdot f_{n+1}$ to w, we obtain $R(w)$ which is recognized by A'. So $R(w) \in \mathscr{L}(A') \subseteq \mathscr{L}(R(A))$.

□

However, we do not have $R(\mathscr{L}(A)) \supseteq \mathscr{L}(R(A))$ as shown in the following example:

Example 2. Let A the lattice automaton that recognizes the language $L = \{[0,0],[1,1]\}$ (i.e. there is only one process, with one integer variable which values is either 0 or 1) and let R be the rewriting rule $\top / f_1.f_2$ where $f_1 : x \mapsto 2x$ and $f_2 : x \mapsto 4x$ then $R(\mathscr{L}(A)) = \{[0,0].[0,0]\,,\,[2,2].[4,4]\}$. But $R(A)$ is a lattice automaton that can recognize 4 words:
$\mathscr{L}(R(A)) = \{[0,0].[0,0]\,,\,[2,2].[4,4]\,,\,[0,0].[4,4]\,,\,[2,2].[0,0]\}$.

Application of a transducer The following algorithm computes the application of a symbolic transducer $T = \langle \Lambda, Q^T, Q_0^T, Q_f^T, \Delta^T \rangle$ to a (language recognized by a) lattice automaton $A = \langle \Lambda, Q^A, Q_0^A, Q_f^A, \Delta^A \rangle$. The idea is to consider the cartesian product $Q^T \times Q^A$ and to create transitions whenever it is allowed by the transducer and the automaton.

ApplyTransducer(T,A):

$\Delta^{T(A)} = \emptyset$
$\forall (p,q) \in Q^T \times Q^A$
$\forall p' \in Q^T, \forall (p,Q,F,p') \in \Delta^T$ with $G \in \Lambda^n$ and $F : \Lambda^n \to \Lambda^*$
$\forall q' \in Q^A$ such that there is a sequence of transitions $q \xrightarrow{w} q'$ in A with $w \in \Lambda^n$
If $G \sqcap w \neq \bot$ then $\Delta^{T(A)} := \Delta^{T(A)} \cup \{((p,q), F(G \sqcap w), (p',q'))\}$
$\forall (p,q) \in Q_0^{T(A)}$ if $p \in Q_0^T$ and $q \in Q_0^A$
$\forall (p,q) \in Q_f^{T(A)}$ if $p \in Q_f^T$ and $q \in Q_f^A$

Note that $\Delta^{T(A)} = \Delta^{T(A)} \cup \{((p,q), F(G \sqcap w), (p',q'))\}$ means that we add not one but a sequence of transitions (introducing fresh new states). So the set of states of the resulting automata $T(A)$ is the union of $Q^T \times Q^A$ and all the fresh states we added. Fig. 6 gives an illustration of an application of a transducer mapping the semantics of the single local instruction of our program $(l_7, \texttt{[x := x + 4]}, l_8)$ on single letter program state set (i.e. only one process). Please note that, for the sake of clarity, we use line numbers as locations. l_7 is the location just before the evaluation of the assignment. l_8 is, thus, after the evaluation and l_9 represents the last location symbolising the end of a process execution. Our transducer application algorithm complexity is $O(|Q_A| \cdot |Q_T| \cdot |\Delta_T| \cdot |\pi|^N)$ where π is the lattice automata's partition (its size depends on the locations of the program) and N is the maximum length of all transition guards (here $N = 1$).

Theorem 2. *Let T be a symbolic transducer and A a lattice automaton. We have: $T(\mathscr{L}(A)) \subseteq \mathscr{L}(T(A))$.*

Proof. Let $T = \langle \Lambda, Q^T, Q_0^T, Q_f^T, \Delta^T \rangle$ and $A = \langle \Lambda, Q^A, Q_0^A, Q_f^A, \Delta^A \rangle$. Let $w \in \mathscr{L}(A)$; we must prove that $T(w) \subseteq \mathscr{L}(T(A))$. By definition, $w' \in T(w)$ if $p_0 \xrightarrow{w_1/w'_1} \twoheadrightarrow$

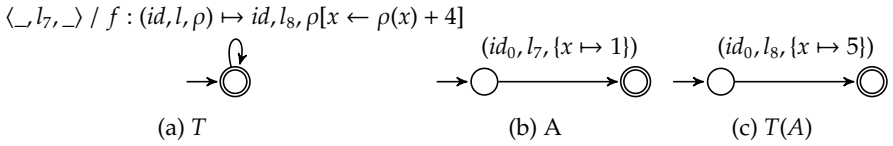

(a) T (b) A (c) $T(A)$

Fig. 6: Transducer application

$p_1 \xrightarrow{w_2/w'2}\!\!\!\!\!\rightarrow \ldots \xrightarrow{w_n/w'_n}\!\!\!\!\!\rightarrow p_n$ with $p_0 \in Q_0^T, p_n \in Q_f^T$, $w = w_1.w_2 \ldots w_n$, $w' = w'_1.w'_2 \ldots w'_n$ and $\forall i = 1..n, \exists (p_{i-1}, G_i, F_i, p_i) \in \Delta^T$ with $w_i \sqsubseteq G_i$ and $w'_i = F(w_i)$.

Since $w \in \mathscr{L}(A)$, it means $\exists q_0 ... q_n \in Q^A$ with $q_0 \in Q_0^A, q_n \in Q_f^A$ such that, $q_0 \xrightarrow{w_1}\!\!\!\!\!\rightarrow q_1 \xrightarrow{w_2}\!\!\!\!\!\rightarrow \ldots \xrightarrow{w_n}\!\!\!\!\!\rightarrow q_n$. In other words $\forall i = 1..n$ there is a sequence of transitions: $q_{i-1} \xrightarrow{\lambda_{i,1}} \ldots \xrightarrow{\lambda_{i,m_i}} q_i$ with $w_i \sqsubseteq \lambda i, 1 \ldots \lambda_{i,m_i}$. So $w_i \subseteq G_i \sqcap \lambda i, 1 \ldots \lambda_{i,m_i}$ and $F(w_i) = w'_i \sqsubseteq F(G_i \sqcap \lambda i, 1 \ldots \lambda_{i,m_i})$. By definition of $T(A)$, $(p_0, q_0) \xrightarrow{w'_1}\!\!\!\!\!\rightarrow \ldots \xrightarrow{w'_n}\!\!\!\!\!\rightarrow (p_n, q_n)$ is an accepting run of $T(A)$, thus $w' \in \mathscr{L}(T(A))$.

The same principle applies for R(A).

□

We note $T_{ext}(A) = R(A) \cup T(A)$ the automaton resulting of the union of `ApplyRule(R,A)` and `ApplyTransducer(T,A)`.

3.4 Fixpoint computation

As we said before, the reachability set is defined as the fixpoint $\text{Post}^*_\tau(I)$; If we can compute $T^*_{ext}(I)$ in the abstract domain of lattice automata, we will get an over-approximation interpretation of this reachability set. However, there are infinitely increasing sequences in this abstract domain, so we need to apply a *widening operator* to ensure the termination of the computation. There exists a widening operator which "lifts" a widening operator ∇_Λ defined for Λ to the abstract domain of lattice automata: $A_1 \nabla A_2$ applies ∇_Λ to each transition of A_1 and A_2 when the two automata have the same "shape"; otherwise, it merges some states of $A_1 \cup A_2$ to obtain an over-approximation (see [9]).

The generic fixpoint algorithm is thus to apply the widening operator ∇ at each step until we reach a post-fixpoint, i.e. we iterate the operator

$$T_\nabla(S) = \begin{cases} S & \text{if } T_{ext}(S) \sqsubseteq S \\ S \nabla (S \cup T_{ext}(S)) & \text{otherwise} \end{cases}$$

This computation gives a post-fixpoint $T^\infty \supseteq T_{ext}(I)$. In practice, this method may yield very imprecise upper bounds. Since Λ contains information about the location of each process, we can improve the precision by applying ∇_Λ only to locations corresponding to an entry point of a loop. It is known [3] that we only

need widening to break dependency cycles and [3] gives an extensive study on the choice of widening application locations.

Once we get an over-approximation of the reachability set, we can check any safety property expressed as a set of bad states represented by a lattice automaton B; if $T^{\infty} \cap B = \emptyset$, then the system is safe. If not, the property may be false, thus we raise an alarm.

On our example (Fig. 1), applying our method using a precise relational numerical abstract domain (e.g. polyhedra) gives us a reachability set. We can prove the safety property given on Fig. 2 by using the following invariant present in the reachability set:

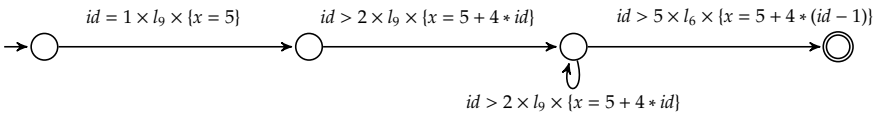

4 Verification of MPI programs

In order to validate our approach, we applied our method to the Message Passing Interface (MPI). MPI is a specification of a message passing model. Many implementations have been developed and it is widely used in parallel computing for designing distributed programs. Every process has its own memory and *shares a common code*. A notion of *rank* (acting as id) is present in order to differentiate the processes. This paradigm makes a good candidate to map our model onto.

We developed a prototype[3] that targets a MPI subset for the C language. It currently supports *synchronous MPI communications*, integer and floating point values as well as a good subset of the C language. Currently, we do not support dynamic process creation in MPI.This prototype has been implemented as a *Frama-C* plug-in.This plug-in uses a lattice transducer library we developed on top of an existing lattice automata implementation. Our abstract domains are given by the `Apron` library. This prototype has been written in OCaml. The current size of the plug-in is around 10.000 lines of code and is still a work in progress. Unfortunately, due to licensing issues, its source code is not available yet.

To illustrate our method, we refer, throughout this section, to a small MPI program (Fig. 7). This program runs N processes that each computes $1/2^{(rank+1)}$. Then, the root (i.e. rank = 0) process collects each local result and sums them by a call to the `MPI_Reduce` primitive.

[3] The prototype can be found at: `https://www-apr.lip6.fr/~botbol/mpai`

```
1   int main(int argc, char **argv) {
2     int rank, i;
3     float res, total;
4     MPI_Init(&argc, &argv);
5     MPI_Comm_rank(MPI_COMM_WORLD, &rank);
6     i = 1 << (rank + 1);
7     res = 1. / i;
8     MPI_Reduce(&res, &total, 1, MPI_FLOAT, MPI_SUM, 0, MPI_COMM_WORLD);
9     MPI_Finalize();
10    return 0;
11  }
```

Fig. 7: MPI program computing: $\sum_{i=1}^{n} \frac{1}{2^i}$

4.1 Program state representation

Each (abstract) local process state is a tuple $\langle l, \lambda \rangle \in L \times \Lambda$, where L is the set of locations and Λ a numerical abstract lattice. In the examples of this section, Λ is the lattice of Intervals. Moreover, we distinguish the value of Id from the other variables.

To illustrate, we give the initial configuration with 2 processes starting at `MPI_Init(&argc, &argv)` (variable declarations are omitted) and represented as a lattice automaton. At this point, each environment variable is set to $\top$ meaning they are not initialized and can have any possible value.

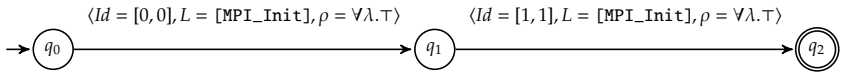

4.2 Transducer automatic generation

Starting from a MPI/C program, the goal is to automatically generate a lattice transducer that fully encodes the program semantics. To achieve that, we first compute the program's Control Flow Graph (CFG). Then, we translate each CFG transition into a lattice transducer rule yielding the complete transducer encoding the program semantics.

As stated before, we differentiate local instructions that affects only one process at a time from global instructions, such as MPI communications, that modify the global state of the program. The translation of local instructions is straight-forward: we use classical transfer functions that are defined in the Apron library to evaluate the expressions and do the assignments. As shown below, an "if" C statement will be translated into two corresponding rules for both condition cases.

```
1  if (x > 10){
2     ...
3  } else {
4     ...
5  }
6        ...
```

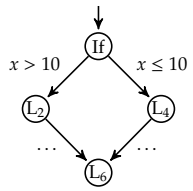

C program CFG

$$\top \times \texttt{[If]} \times \{x \in [11, +\infty]\} \;/\; f : (id, L, \rho) \mapsto id, [\mathrm{L}_2], \rho$$

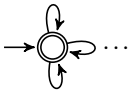

$$\top \times \texttt{[If]} \times \{x \in [-\infty, 10]\} \;/\; f : (id, L, \rho) \mapsto id, [\mathrm{L}_4], \rho$$

Resulting transducer

Below is the transducer generated from all local instructions of the MPI program depicted on Fig. 7. Note that, with this set of local rules, there is no way to evolve from the `MPI_Reduce` location. As mentioned in the previous section, we dissociate the global rules from the transducer's local rules. Therefore, this transition will be presented in the next section. Finally, in order to model process inactivity, we add a simple rule $\top \;/\; f : x \mapsto x$ meaning that any process at any location might not evolve.

$$\top \times \texttt{[MPI\_Init]} \times \top \;/\; f : (id, l, \rho) \mapsto id, \texttt{[MPI\_Comm\_rank]}, \rho$$
$$\top \times \texttt{[MPI\_Comm\_rank]} \times \top \;/\; f : (id, l, \rho) \mapsto id, \texttt{[i = 1 << (rank + 1)]}, \rho[rank \leftarrow id]$$
$$\top \times \texttt{[i = 1 << (rank + 1)]} \times \top \;/\; f : (id, l, \rho) \mapsto id, \texttt{[res = 1. / i]}, \rho[i \leftarrow 1 << (\rho(rank) + 1)]$$

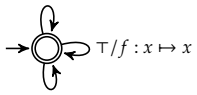

$$\top \times \texttt{[res = 1. / i]} \times \top \;/\; f : (id, l, \rho) \mapsto id, \texttt{[MPI\_Reduce]}, \rho[res \leftarrow 1./\rho(i)]$$
$$\top \times \texttt{[MPI\_Finalize]} \times \top \;/\; f : (id, l, \rho) \mapsto id, \texttt{[return]}, \rho$$

4.3 Encoding communication primitives

Our prototype currently accepts this subset of MPI primitives : `MPI_Send`, `MPI_Recv`, `MPI_Bcast`, `MPI_Comm_rank`, `MPI_Comm_size` and `MPI_Reduce`. We already described the symbolic rewriting rules in Sec. 3 except for `MPI_Comm_rank`, `MPI_Comm_size`, which returns the id of the current process and the total number of processes, and `MPI_Reduce`. Let us give the semantics of the last one:

```
MPI_Reduce(void* send_data, void* recv_data, int count,
      MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm communicator)
```

This global communication primitive gathers every process' `send_data` buffer and applies a commutative (the order of reduction is undefined) op-

erator `op` between every value. The result is then sent to the process of rank `root` at its `recv_data` address. `count` and `datatype` are respectively the size of these buffers and the type of each value. The `communicator` defines a group of processes where the communication will occur. We assume a single group.

We cannot represent this global communication in our model with only one rule. Our solution is to break it down into three different ones. The main idea is to spawn a "collector" process that will be in charge of gathering each process' `send_data` and applying the reduction operation on its accumulator. This collector will move through the program state (i.e. a word of local states) by swapping, at each iteration, with the next process. Before starting to move this collector, we have to ensure that no involved process might evolve. Therefore, we *lock* them using a special location. When the collector reaches the end of the word, it sends its accumulator to the root process through a point-to-point communication, destroys itself and, finally, unlocks the processes. The three skeleton rules used in our prototype are given here:

1. $G_1 = (\top \times \texttt{[MPI\_Reduce]} \times \top)^*$
   $F_1 = (f : \_ \mapsto -1, \texttt{[Collector]}, \{recv\_data \mapsto e\}) \cdot$
   $(F : ((id, l, \rho), \_) \mapsto id, \mathrm{lock}(l), \rho)$ where e is the neutral element of `op`
2. $G_2 = (\top \times \texttt{[MPI\_Reduce]}_{\mathrm{lock}} \times \top)^* \cdot$
   $(\top \times \texttt{[Collector]} \times \top) \cdot (\top \times \texttt{[MPI\_Reduce]}_{\mathrm{lock}} \times \top) \cdot$
   $(\top \times \texttt{[MPI\_Reduce]}_{\mathrm{lock}} \times \top)^*$
   $F_2 = \texttt{Id}^* \cdot$
   $(f : ((id_{coll}, l_{coll}, \rho_{coll}), (id_{proc}, l_{proc}, \rho_{proc})) \mapsto id_{proc}, l_{proc}, \rho_{proc}) \cdot$
   $(f : ((id_{coll}, l_{coll}, \rho_{coll}), (id_{proc}, l_{proc}, \rho_{proc})) \mapsto$
   $id_{coll}, l_{coll}, \rho_{coll}[\texttt{recv\_data} \leftarrow \rho_{coll}(\texttt{recv\_data}) \, [\![\texttt{op}]\!] \, \rho_{proc}(\texttt{send\_data})]) \cdot$
   $\texttt{Id}^*$
3. $G_3 = (\top \times \texttt{[MPI\_Reduce]}_{\mathrm{lock}} \times \top)^* \cdot (\texttt{root} \times \texttt{[MPI\_Reduce]}_{\mathrm{lock}} \times \top) \cdot$
   $(\top \times \texttt{[MPI\_Reduce]}_{\mathrm{lock}} \times \top)^* \cdot (\top \times \texttt{[Collector]} \times \top)$
   $F_3 = (F : ((id, l, \rho), \_) \mapsto id, \texttt{[next\_loc]}, \rho)$
   $(f : ((id_{root}, l_{root}, \rho_{root}), (id_{coll}, l_{coll}, \rho_{coll})) \mapsto$
   $id_{root}, \texttt{[next\_loc]}, \rho_{proc}[\texttt{recv\_data} \leftarrow \rho_{coll}(\mathrm{recv\_data})]) \cdot$
   $(F : ((id, l, \rho), \_) \mapsto id, \texttt{[next\_loc]}, \rho)$

For our example program, our prototype automatically instantiates these rules in a set R. Fig. 8, illustrates the iterative applications of R on program state A where both processes have reached the `MPI_Reduce` location by successive application of the transducer T on the initial configuration I.

We give in Fig. 9 the full reachability set computed by our prototype. For readability purposes, we do not show variables that are not set (i.e. when their value is $\top$). We also factorize the transitions: for each multiple transitions from

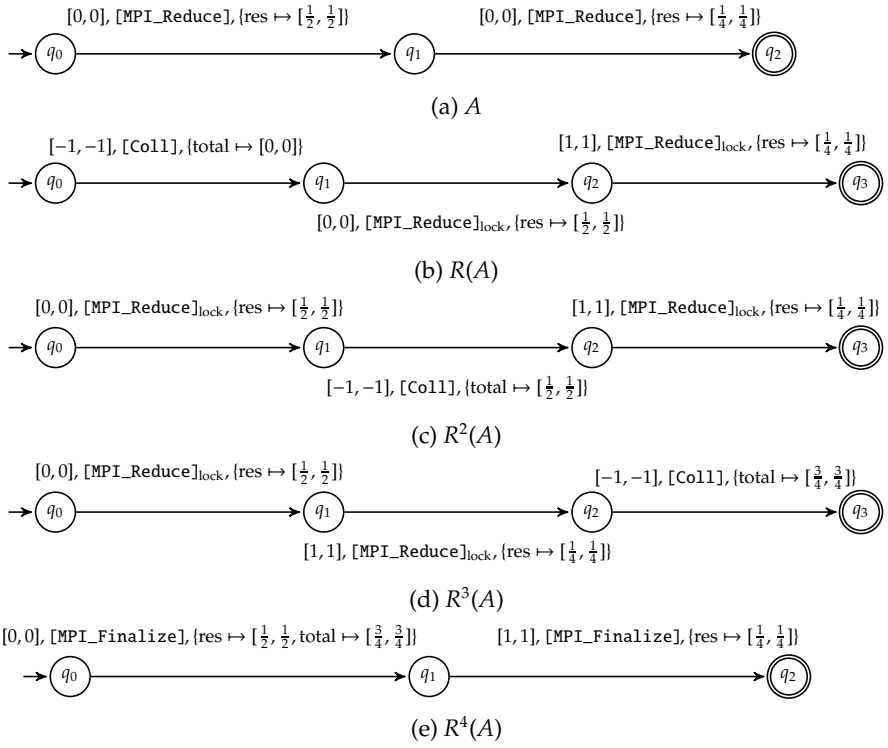
[0, 0], [MPI_Reduce], {res ↦ [1/2, 1/2]}
[0, 0], [MPI_Reduce], {res ↦ [1/4, 1/4]}
q0
q1
q2
(a) A
[−1, −1], [Coll], {total ↦ [0, 0]}
[1, 1], [MPI_Reduce]lock, {res ↦ [1/4, 1/4]}
[0, 0], [MPI_Reduce]lock, {res ↦ [1/2, 1/2]}
q3
(b) R(A)
[0, 0], [MPI_Reduce]lock, {res ↦ [1/2, 1/2]}
[1, 1], [MPI_Reduce]lock, {res ↦ [1/4, 1/4]}
[−1, −1], [Coll], {total ↦ [1/2, 1/2]}
(c) R²(A)
[0, 0], [MPI_Reduce]lock, {res ↦ [1/2, 1/2]}
[−1, −1], [Coll], {total ↦ [3/4, 3/4]}
[1, 1], [MPI_Reduce]lock, {res ↦ [1/4, 1/4]}
(d) R³(A)
[0, 0], [MPI_Finalize], {res ↦ [1/2, 1/2, total ↦ [3/4, 3/4]}
[1, 1], [MPI_Finalize], {res ↦ [1/4, 1/4]}
(e) R⁴(A)

Fig. 8: Application of R on a configuration

a node p to q, we merge them into a single one and concatenate their labelled local process states (i.e. $(p, \sigma, q); (p, \sigma', q); \Rightarrow (p, \sigma; \sigma', q)$).

5 Experiments

We present in this section some of the analysis results of our prototype. We found several tools that provide a formal verification of MPI programs. One of the most advanced we found is called "In-situ Partial Order" [16] (ISP). It is based on model checking and performs a dynamic analysis in order to detect the presence of *deadlocks*. To the best of our knowledge, our tool is the only one that computes the reachability set. We present some examples where we verify numerical properties and although our prototype focuses on safety properties, it can also detects that program states (i.e. words) in our set are not matched by any rules and therefore detect deadlocks. In these cases, we can raise an alarm (which can be false ones due to our abstractions).

Program	LoC	nb proc.	state space size	nb nodes	nb transitions	exec. time
deadlock random	23	2	225	4	20	0.5s
dining philosophers	42	4	83521	17	112	8s
dining philosophers	42	6	24137569	54	447	1232s
sum program	11	50	∞	200	601	86s
pi approximation	26	50	∞	200	751	104s

We tested our prototype on several examples. We display here the results of significant ones. Our parameters are: the number of processes we start with, the concrete state space size, the number of nodes in the lattice automaton that represents the final reachability set, its number of transitions and the execution time. The concrete state space size is the enumeration of all possible program states; it is infinite when there are integer variables. We prove on these examples two kinds of properties: deadlock detection and numerical safety properties.

First is a potentially deadlocking program "random deadlock" where two process tries to communicates randomly: both test a random condition that leads respectively to a send or a receive call towards the other process. As ISP is dynamic and depends on the MPI execution, it will not always detect this simple deadlock. However, as we compute the reachability set, we easily observe this deadlock and can raise an alarm.

We implemented a MPI version of the dining philosopher problem where philosophers and forks are processes. The forks processes will give permission to "pick them up" and "put them down" modeled by point-to-point communications. Naturally, the program has deadlocks and again the reachability set exhibits them. The growth in computation time is explained by the amount of possible interleavings that our algorithm is currently not capable of filtering and by the precision we wish to attain (thus, no strong abstractions) in order to precisely determine the deadlocks (and not a false alarm).

The next two following examples both implement a floating point value approximation. The first one is our example program used in Fig. 7. The same property is used: total $\in [0, 1]$. The second one is a computation on pi based on the approximation of $\int_0^1 \frac{4}{1+x^2}$ with sums of n intervals dispatched on n processes. Again the property is a framing of the result ($\in [3, 4]$) . These two examples display the capacity of our prototype to handle real-life computations. However, we would like to generalize these two examples to any number of process. We can model an initial configuration with an unbounded number of process and run our analysis on it. Unfortunately, we cannot infer a relation between the process rank or the number of processes with our current numerical domains. Therefore, our *sum program*'s analysis, on an unbounded number of process, can detect that each process computes a local result $\in [0, \frac{1}{2}]$ but the total sum will be abstracted to $[0, +\infty]$.

6 Conclusion

We presented a new way to do static analysis on a model of concurrent programs that allow unicast and multicast communication as well as dynamic process creation. We described the general framework of the method with well-founded abstraction of the semantics and program states. We applied our technique in order to compute reachability sets of MPI concurrent programs with numerical abstract domains. We showed that building such an analysis on a realistic language, such as MPI/C programs, is feasible and yields encouraging results. Moreover, abstract interpretation allows us to verify numerical properties which was not done before on such programs, and the lattice automata allow the analysis to represent (and automatically discover) regular invariants on the whole program states.

Future work includes theoretical and practical improvements of our analyzer, especially the application algorithm which is currently not optimized. One way to do that is to run a quick pre-analysis using a simple, non-numerical abstract domain to obtain information (e.g. rewriting rules that are never activated), so that we may simplify the rules before using more costly numerical abstract domains. We also wish to design a specification language allowing us to write regular properties more easily. We will also improve our analyser by taking into account more MPI primitives as well as supporting general C constructs (pointers, functions, etc.) thanks to better interactions with the other Frama-C plug-ins. Finally, we will deal with asynchronous communications (FIFO queues) and shared variables using non-standard semantics and/or a reduced product with abstract domains that can efficiently abstract these kind of data.

References

1. Parosh Aziz Abdulla, Bengt Jonsson, Marcus Nilsson, and Mayank Saksena. A Survey of Regular Model Checking. In *CONCUR 2004 - Concurrency Theory, 15th International Conference, London, UK, 2004, Proceedings*, pages 35–48, 2004.
2. Ahmed Bouajjani, Peter Habermehl, and Tomás Vojnar. Abstract Regular Model Checking. In *Computer Aided Verification, 16th International Conference, CAV 2004, Boston, MA, USA, July 13-17, 2004, Proceedings*, pages 372–386, 2004.
3. François Bourdoncle. Efficient Chaotic Iteration Strategies With Widenings. In *Proceedings of the International Conference on Formal Methods in Programming and their Applications*, pages 128–141. Springer-Verlag, 1993.
4. Jean-Loup Carre and Charles Hymans. From Single-thread to Multithreaded: An Efficient Static Analysis Algorithm. *ArXiv e-prints*, October 2009.
5. Patrick Cousot and Radhia Cousot. Abstract Interpretation: A Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fixpoints. In *Conference Record of POPL'77 : The 4th ACM Symposium on Principles of Programming Languages, 1977*, pages 238–252, 1977.

6. Patrick Cousot and Nicolas Halbwachs. Automatic Discovery of Linear Restraints Among Variables of a Program. In *Conference Record of POPL'78 : The 5th ACM Symposium on Principles of Programming Languages*, POPL '78, pages 84–96, 1978.
7. Jérôme Feret. Partitioning the Threads of a Mobile System. *CoRR*, abs/0802.0188, 2008.
8. Pietro Ferrara. Checkmate: A Generic Static Analyzer of Java Multithreaded Programs. In *2009 Seventh IEEE International Conference on Software Engineering and Formal Methods*, pages 169–178, Nov 2009.
9. Tristan Le Gall and Bertrand Jeannet. Lattice Automata: A Representation for Languages on Infinite Alphabets, and Some Applications to Verification. In *Static Analysis, 14th International Symposium, SAS 2007*, pages 52–68, 2007.
10. Bertrand Jeannet and Antoine Miné. APRON: A Library of Numerical Abstract Domains for Static Analysis. In Ahmed Bouajjani and Oded Maler, editors, *CAV'2009 - Computer Aided Verification*, volume 5643 of *LNCS*, pages 661–667. Springer, 2009.
11. Florent Kirchner, Nikolai Kosmatov, Virgile Prevosto, Julien Signoles, and Boris Yakobowski. Frama-C: A Software Analysis Perspective. *Formal Aspects of Computing*, 27(3):573–609, 2015.
12. Antoine Miné. The Octagon Abstract Domain. *Higher Order Symbol. Comput.*, 19(1):31–100, March 2006.
13. Antoine Miné. Relational Thread-Modular Static Value Analysis by Abstract Interpretation. In *Verification, Model Checking, and Abstract Interpretation - 15th International Conference, VMCAI*, pages 39–58, 2014.
14. Marc Snir, Steve Otto, Steven Huss-Lederman, David Walker, and Jack Dongarra. *MPI-The Complete Reference, Volume 1: The MPI Core*. MIT Press, Cambridge, MA, USA, 2nd. (revised) edition, 1998.
15. Margus Veanes, Pieter Hooimeijer, Benjamin Livshits, David Molnar, and Nikolaj Bjørner. Symbolic finite state transducers: algorithms and applications. In *Conference Record of POPL'12 : The 39th ACM Symposium on Principles of Programming Languages*, pages 137–150, 2012.
16. Anh Vo, Sarvani Vakkalanka, Michael DeLisi, Ganesh Gopalakrishnan, Robert M Kirby, and Rajeev Thakur. Formal verification of practical MPI programs. *ACM Sigplan Notices*, 44(4):261–270, 2009.

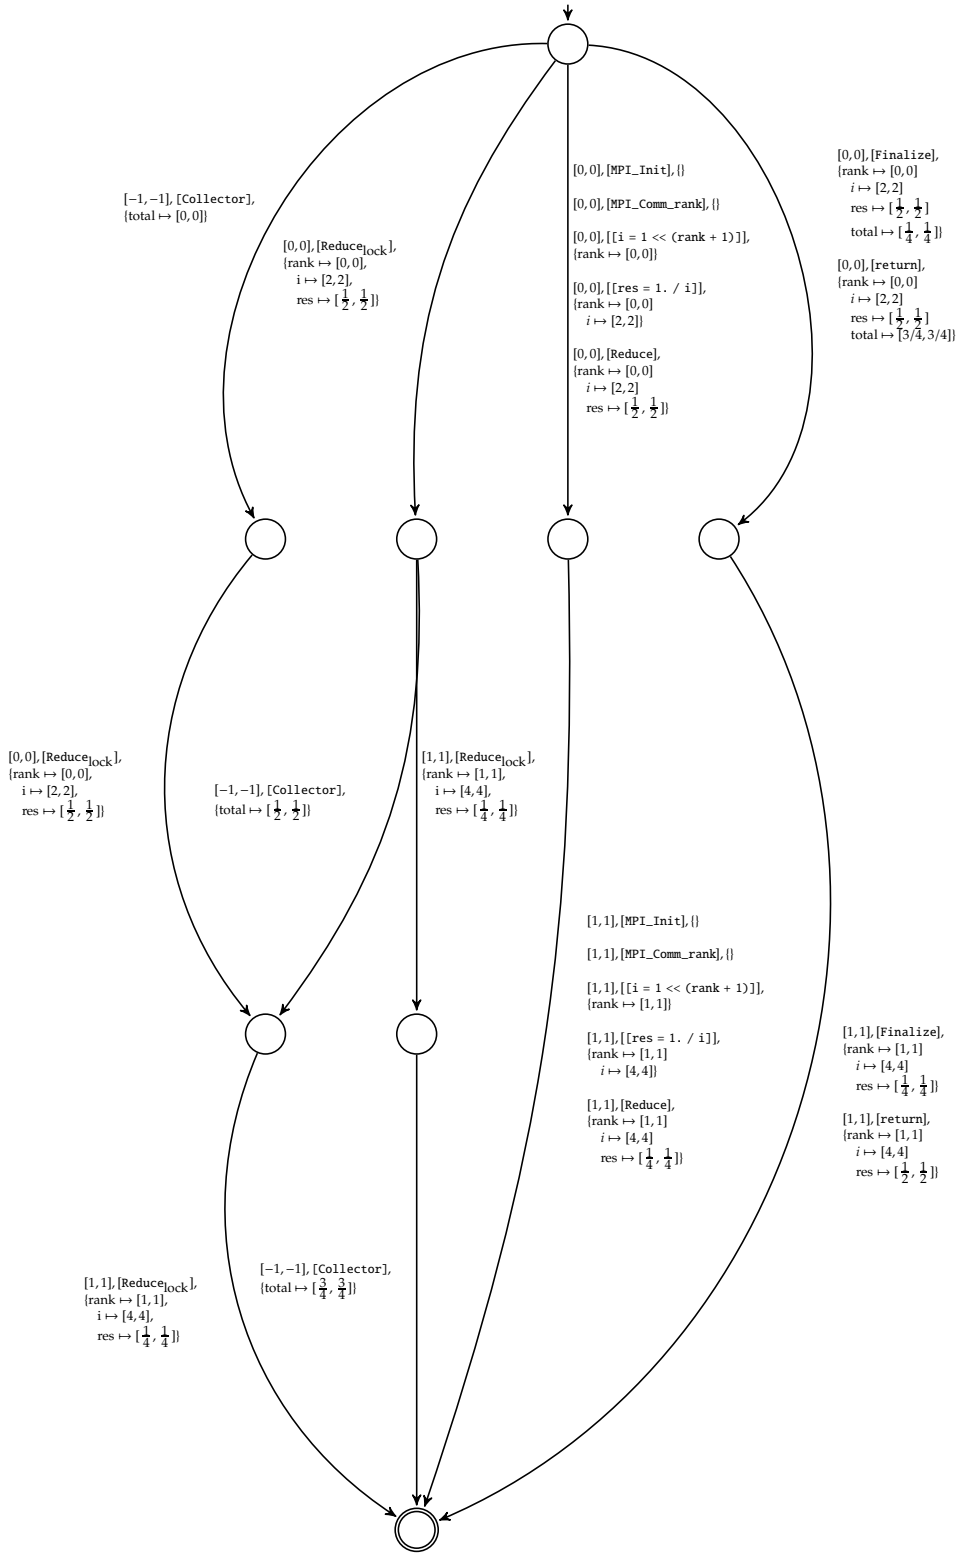
[−1, −1], [Collector],
{total ↦ [0, 0]}
[0, 0], [Reduce_lock],
{rank ↦ [0, 0],
i ↦ [2, 2],
res ↦ [1/2, 1/2]}
[0, 0], [MPI_Init], {}
[0, 0], [MPI_Comm_rank], {}
[0, 0], [[i = 1 << (rank + 1)]],
{rank ↦ [0, 0]}
[0, 0], [[res = 1. / i]],
{rank ↦ [0, 0]
i ↦ [2, 2]}
[0, 0], [Reduce],
{rank ↦ [0, 0]
i ↦ [2, 2]
res ↦ [1/2, 1/2]}
[0, 0], [Finalize],
{rank ↦ [0, 0]
i ↦ [2, 2]
res ↦ [1/2, 1/2]
total ↦ [1/4, 1/4]}
[0, 0], [return],
{rank ↦ [0, 0]
i ↦ [2, 2]
res ↦ [1/2, 1/2]
total ↦ [3/4, 3/4]}
[0, 0], [Reduce_lock],
{rank ↦ [0, 0],
i ↦ [2, 2],
res ↦ [1/2, 1/2]}
[−1, −1], [Collector],
{total ↦ [1/2, 1/2]}
[1, 1], [Reduce_lock],
{rank ↦ [1, 1],
i ↦ [4, 4],
res ↦ [1/4, 1/4]}
[1, 1], [MPI_Init], {}
[1, 1], [MPI_Comm_rank], {}
[1, 1], [[i = 1 << (rank + 1)]],
{rank ↦ [1, 1]}
[1, 1], [[res = 1. / i]],
{rank ↦ [1, 1]
i ↦ [4, 4]}
[1, 1], [Reduce],
{rank ↦ [1, 1]
i ↦ [4, 4]
res ↦ [1/4, 1/4]}
[1, 1], [Finalize],
{rank ↦ [1, 1]
i ↦ [4, 4]
res ↦ [1/4, 1/4]}
[1, 1], [return],
{rank ↦ [1, 1]
i ↦ [4, 4]
res ↦ [1/2, 1/2]}
[1, 1], [Reduce_lock],
{rank ↦ [1, 1],
i ↦ [4, 4],
res ↦ [1/4, 1/4]}
[−1, −1], [Collector],
{total ↦ [3/4, 3/4]}

Fig. 9: Sum program reachability set with 2 processes